

Report of the Task Force on Experimental Protection, July 2015

Members: James Dunlop (chair), Michael Blaskiewicz, Mickey Chiu, Bill Christie, Angelika Drees, Wolfram Fischer, Xiaofeng Gu, Chuyu Liu, Guillaume Robert-Demolaize

Executive Summary

This task force has been charged with assessing the risk of damage to the detectors of an abort kicker prefire within the proposed running conditions in Run 16, given past experience, along with assessing possible ways to ameliorate this risk. The risk of damage to the detectors is the product of the probability of a prefire occurring in any machine configuration and the severity of damage should a prefire occur. There is currently some protection in the machine for damage to the experiments due to a prefire, but this was shown to be inadequate in Run 15 with the damage to the PHENIX Muon Piston Calorimeter (MPC) and Silicon Vertex (VTX) detectors during the p+Au running at $\sqrt{s_{NN}} = 200$ GeV. The working hypothesis is that this protection is inadequate because of the closeness of the masks to the PHENIX intersection region in conjunction with the reduced beam aperture at the DX magnets due to the different charge-to-mass ratios of the p and Au beams. This results in showering of secondary particles into the detector when a prefire event occurs and primary particles are lost at DX magnet in front of the PHENIX detector. The protection was adequate for configurations with beams of the same charge-to-mass ratio in both rings in Run-13 and Run-14. The STAR MTD, which sustained damage in previous runs due to a prefire, sustained no damage in Run 15, presumably due to the orbit bumps and masks, installed in response to the initial damage, that deposit the beam far from the detector.

Based on historical data, the rate of abort kicker pre-fires is seen to depend on the voltage at which the abort kicker power supplies operate. A detailed analysis of the historical data leads to the following estimate of the probabilities for a run plan optimized to put the largest risk latest, and so maximize the chance that the bulk of the data be collected without incident:

p+Au @ 20 GeV for 2 weeks, with total pre-fire probability: (0.54 +/- 0.54)%

p+Au @ 39 GeV for 2 weeks, with total pre-fire probability: (0.95 +/- 0.95)%

p+Au @ 62 GeV for 1.5 weeks, with total pre-fire probability: (1.4 +/- 1.4)%

Under this scenario, the chance is 97+/-3% that the run would occur without incident.

The PHENIX MPC was severely damaged by abort kicker prefires under two running conditions, in Run 12 Cu+Au and Run 15 p+Au. In both cases, key transistors in the MPC pre-amplifier electronics sustained overcurrent damage. Repairs of these circuits are possible during a summer shutdown, but the repairs cannot be done during a run. Bench studies have been able to reproduce this failure mode with signal size equivalent to approximately 30 TeV of energy (i.e. the energy of 1.5 Au ions at 100 GeV/nucleon) deposited into a single MPC crystal. A protection circuit has been designed and tested, which will raise this protection threshold by a factor of 400 and may completely protect the component from failure. This protection circuit will be installed on the MPC electronics during the shutdown prior to Run 16.

No further protection will be available for the VTX stripixel detector, which sustained irreparable damage in Run 15 over a significant fraction of the detector. The PHENIX

collaboration has stated that the physics program of the low-energy p+Au running does not critically depend on the use of this detector, and there are no plans to re-use the detector after Run 16.

Studies of the instantaneous radiation dose into the PHENIX detector from an abort kicker prefire, using particle tracking and GEANT, have begun, but could not be completed to provide a quantitative answer in the timescale of this task force. Therefore the task force cannot quantify the effect of changes in the machine conditions on the instantaneous dose into PHENIX, should a prefire occur. This includes the preferability of Au+p over p+Au running and the effect of orbit changes. That said, major damage has only occurred in asymmetric running, in which the orbits of the beams come close to the DX magnets. It was realized during the task force meetings that the DX magnets could be moved for d+Au running to make the apertures nearly equivalent to that of symmetric running, during which there has been only minor damage to the MPC (twice) despite several more prefires. At lower energies there is also enough margin in power supplies and magnets to maximize the aperture for the beam coming into PHENIX at the DX magnet as long as there is sufficient aperture for the outgoing beam. Further optimization of aperture at the DX may be possible by optimizing the location of the DX magnets, especially in the Au+p configuration that has never been run, but this needs investigation into the limitations imposed by shielding configurations.

There are further measures that can be taken to reduce the chance of prefires. One possibility would be to install, in series with the existing gas discharge thyratrons, mechanical relays capable of switching state in several milliseconds. This project is estimated to cost ~\$300k to \$400k, and would necessitate a detailed machine safety review, since the mechanical relays increase the latency between an abort trigger and the firing of the abort kickers. A further measure would be to move the abort dumps and masks to another place in the machine, well away from PHENIX. A compact movable mask (\$50-70k) after the beam dump was proposed for detector protection as well. It reduces possible pre-fire beam loss by 1-2 orders of magnitude for any species and any existing lattice design. This mask can be used as protection not only for PHENIX detectors, but also for STAR detectors and CeC wiggler. These options are not possible in the summer shutdown before Run 16, due to the scale of the work, but should be actively investigated prior to running with sPHENIX.

Probability Assessment of Prefires at Low Energies

The RHIC beam abort kicker system consists of five identical systems for each beam, each comprising a pulsed power supply driving a kicker magnet [1]. The pulsed power supplies deliver a 12.8 μ s long pulse, with a peak voltage and current of 27 kV and 21 kA at the maximum RHIC energies of $\sqrt{s} = 200$ GeV Au+Au of 510 GeV p+p. Each power supply consists of energy storage capacitors that form a pulse forming network, a thyatron switch to hold off the high voltage until triggered, and the associated control systems. The pulse forming network is designed to modulate the voltage so that when the kicker is activated the magnet is supplied with a varying current, thus spreading the beam in the dump to prevent overheating in a local area.

Prefires have been attributed to various sources, but the leading hypothesis is an unintentional discharge of the deuterium gas thyatron. In some cases, faulty capacitors may have discharged accidentally [2]. Extensive studies have been done of the prefire rates in previous runs by C-AD [2-3]. In the most recent study, the authors report that since 2007, there have been 134 prefires at or near the maximum RHIC energies, one prefire at 200 GeV p+p, and none at low energies [3]. The dependence on abort kicker voltage for the yellow ring is shown in Figure 1. The point at 11 kV corresponds to 200 GeV p+p.

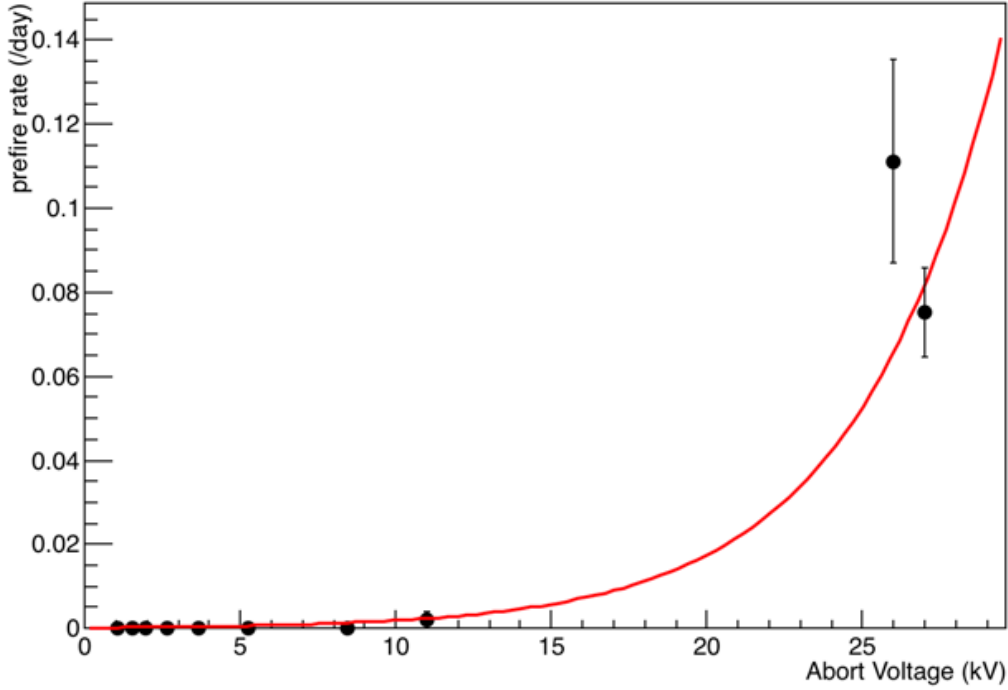


Figure 1: Prefire rates in the yellow ring vs. abort voltage, with a fit to an exponential.

While there seems to be a clear dependence on voltage, one can quantitatively check whether the behavior at low energy is different from that at high energy, i.e., were we just “lucky” to not have any prefires at low energy? We assume that the number of prefires follows a Poisson distribution,

where n = the number of prefires and λ = the mean number of prefires for a given time interval. We also tabulate the amount of time spent at the low energy heavy ion runs and p+p runs in Table 1, since they were not included Ref. [3]. All the runs listed in Ref. [3] are from 2007 and later, since the accounting on prefires was changed then and it is difficult to go back further. In Table 1 we also include the p+p runs from 2002, since it is known that there were no prefires in 200 GeV p+p since then.

System	Energy	Abort Voltage	Weeks running	Days Running
Au+Au	3.85	1.0395	4.6	32.2
Au+Au	5.75	1.5525	1.4	9.8
Au+Au	7.3	1.971	3.4	23.8
Au+Au	9.8	2.646	1.4	9.8
Au+Au	13.5	3.645	1.1	7.7
Au+Au	19.5	5.265	1.8	12.6
Au+Au	31.2	8.424	2.9	20.3
p+p	100	11.0	28.6	200.2
p+p (from 2002)	100	11.0	74.5	521.5

Table 1: Abort voltages and run lengths for 200 GeV p+p energies and below.

Going back to the question of whether low energy running is different from high energy running from the perspective of the abort kickers, we find from Ref. [3] that there was an average rate of 134 prefires over 841 days of running at high energy. This is an average rate of 0.16 prefires per day in RHIC high energy running. The probability that we would have had no prefires in the 113 days of low energy running, assuming that the abort kickers prefire at the same rate as at high energy, is 1.4×10^{-8} . Thus, it seems extremely unlikely that there were no prefires at low energy just due to chance. In fact, based on our understanding of the thyratrons, which are gas-filled tubes which must be ionized to start the arc discharge, we would expect that running them at lower voltages makes them more stable. The same consideration is true for the storage capacitors.

Since it is natural to expect a drop of the prefire rate as one goes to lower voltages, for this study we fit the prefire rate data to an exponential distribution, which decreases sharply with reduced voltage. Of course it could also be that the prefire rate stays constant at the 200 GeV p+p rate when going to lower abort voltages. It may also be that the prefire rate goes to zero at the lower beam energies, if the voltage falls below some critical threshold for the kicker system. The uncertainty in the behavior of the prefire rate at low voltages (<11 kV) is bracketed by the above two scenarios: an upper limit where it is assumed that the prefire rate stays constant at the 200 GeV p+p rate, and a lower limit of zero. The probabilities per week of experiencing a prefire, assuming that the prefire rate follows the exponential given in Figure 1, is

Probability of prefire for 1 week running @ 62.4 = 0.0095
 Probability of prefire for 1 week running @ 39.0 = 0.0047
 Probability of prefire for 1 week running @ 20.0 = 0.0027

For the upper limit the probabilities are roughly twice the above values, and for the lower limit it is zero.

[1] “The RHIC Beam Abort Kicker System”, H. Hahn et al., Proceedings of the 1999 Particle Accelerator Conference, New York

[2] “Analysis of RHIC Beam Dump Pre-fires”, Hahn, H. et al. Conf.Proc. C110328 (2011) 2327-2329 PAC-2011-THP108

[3] “RHIC abort kicker prefire report”epo abort kPerlstein, S., BNL-105539-2014-IR

[4] A. Drees et al., "RHIC prefire protection masks", BNL-107-380-2015-IR

MPC Repair and Protection Modification

The MPC is an electromagnetic calorimeter composed of PbWO₄ crystals with avalanche photodiodes (APD) to read out the scintillation light generated when particles shower in the calorimeter. The APD is readout with a charge-sensitive preamplifier. One crystal assembly is shown in fig. 1. The design is a copy of the Alice PHOS detector, and in particular the preamp was designed to be extraordinarily sensitive, with an equivalent noise charge (ENC) of just 500 electrons, since one of the major PHOS goals was to measure thermal photons.

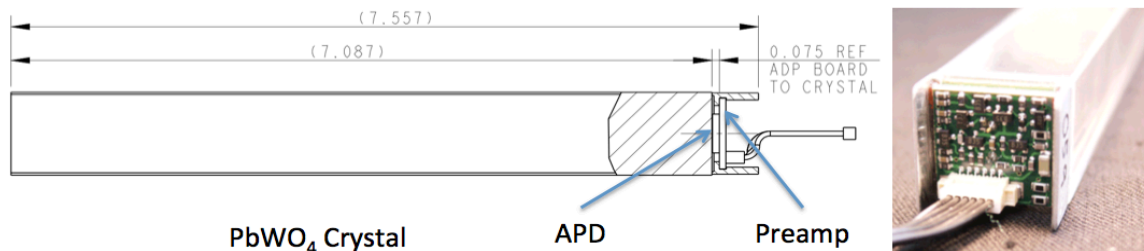


Figure 1: MPC Crystal Assembly, consisting of the PbWO₄, APD glued on the front face, and the preamp with cable attached to the APD and ensclosed in a holder.

The MPC’s location in the forward region of PHENIX, about 8 cm from the beam-pipe places it in a very advantageous location for low-x physics measurements. However, this proximity to the beam-pipe also means that there is the possibility of damage when beam remnants deposit more instantaneous energy than the preamp was designed to handle. This can happen, for instance, when the collider loses control of the beam during an abort kicker prefire event, and up to 10 bunches from the yellow ring are lost near the PHENIX. In the only three instances when major damage occurred to the MPC, all have been associated with prefire events during asymmetric running (2012 Cu+Au and twice during 2015 p+Au). Two of the events have been traced back and determined to have originated from spray coming from the vicinity of the upstream (north)

PHENIX DX magnet. It should be noted that in addition to the three major incidents, there was one incident when about 10 crystals were lost during a Au+Au run.

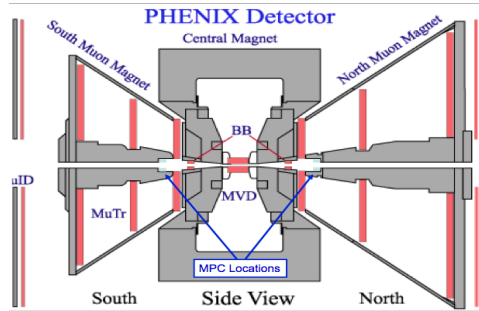


Figure 2: Locations of the MPC in the PHENIX detector

For every MeV of energy deposited the PbWO₄ crystals produces about 1.45 photoelectrons in the APD over a typical scintillation timescale of 30 ns. At a typical APD gain of 50, a current of 0.3 μ A per GeV deposited is expected going into the preamp. At the initial stage of the preamp is a JFET with a 10 mA current limit, which we have verified in the lab holds for even short bursts of current. The JFET is shown in the schematic in fig. 3 below. Energy deposits of about 33 TeV would be enough to damage this event. Typically, we expect a maximum of only about 200 GeV of energy deposited during normal running, so the preamp was designed to handle up to 100x the current enhanced RHIC luminosity.

However, during an abort kicker prefire event, a large amount of energy can and has been deposited in the PbWO₄ scintillating crystal. This large burst of light in the MPC from the spray of the ion bunches hitting the PHENIX IR near the MPC overwhelms the 10 mA current limit, damaging the JFET. The signal comes correlated over a very short period of time since the ion bunches have a RMS width of only ~ 5 ns. A gold ion in RHIC has an energy of 20 TeV energy, not far from the 33 TeV limit of the preamp. However, a gold ion striking the MPC would deposit roughly just 10% of its energy in the struck crystal since the showering process will spread the energy deposition across neighboring crystals. Thus, we expect somewhere around 10 direct gold ions would be enough to destroy one crystal. Since we believe many of the previous events come from the fragments of an upstream interaction, of degraded energy and charge, probably many more particles are needed to destroy the JFET. Geant4 studies are underway to try to quantify this and understand more completely the level of energy deposited during these prefire events.

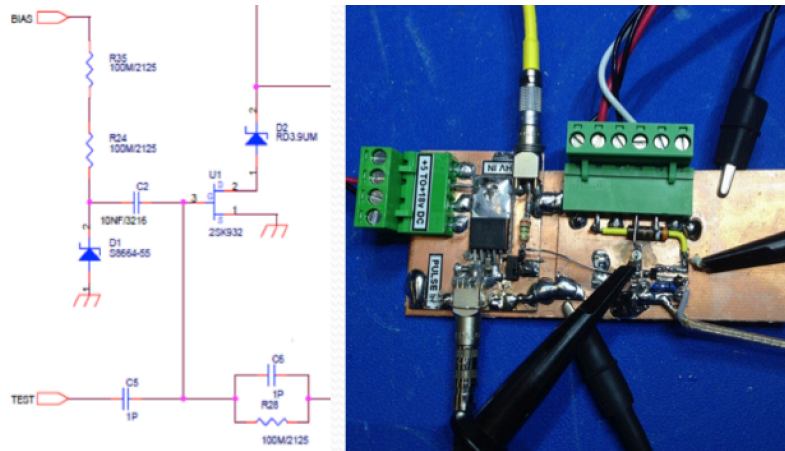


Figure 3: The schematic for the initial stage of the MPC charge sensitive preamp. The APD is D1 and the JFET is U1. The charge integrating circuit can be seen in the lower right of the schematic. The right picture is the test circuit used by Steve Boose.

To understand better the functioning of the MPC preamp, Steve Boose built a test circuit of the initial stage in the preamp, shown in the right of fig. 3. From this it was verified that when the JFET reaches 10 mA or 33 V it would succumb to damage. Previously, in 2012, a study was initiated to test adding a protection diode after the last major MPC damaged incident. We found that various diodes which we expected to work instead made the preamp completely non-functional. From that experience, this time in 2015 we suspected that the previous failures were due to the very tiny leakage current of even just μA or even nA draining the charge integrating capacitor, since this preamp was designed to be extraordinarily sensitive. A BAS416 diode was found for trial as part of a protection circuit. This diode has among the lowest leakage currents available, 3 pA, and is encapsulated in a SOD323 package which makes it possible to install in the cramped preamp space. This time around during testing with LED the preamp remained functional with the diode installed. It should be noted that there are maybe 3-4 diodes with the right parameters to have worked in this preamp.

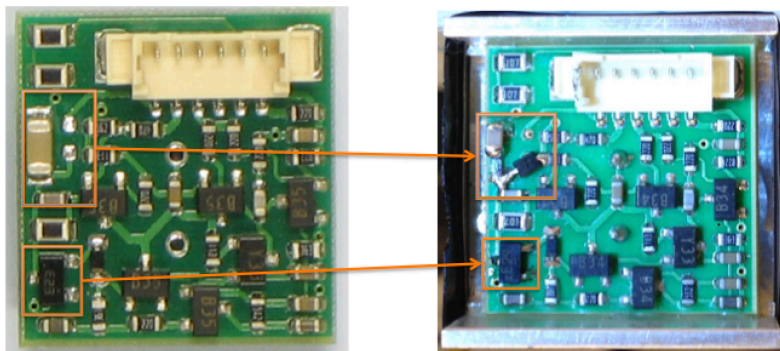


Figure 4: Original preamp and modified preamp. The lower orange boxes indicate the location of the replaced JFET, while the upper box shows the new protection circuit consisting of a diode in series with a resistor, and the replacement capacitor.

To fit this new diode in the space available on the preamp, the input blocking capacitor was replaced with a smaller capacitor. In addition, in order to protect the diode, which can handle a maximum current of 4A, a 100 Ω resistor was added in series (see fig. 4 for a picture of the modified circuit). With this new diode protection system, based on pulse injection tests we have done in the lab we believe the MPC preamps can now sustain at least 400 times the previous level of current with no resistor in series. With a 100 Ω resistor, the MPC preamp should be able to sustain any amount of current, but at the cost of a small increase in noise.